3D Morphable Model Construction for Robust Ear and Face Recognition

John D. Bustard and Mark S. Nixon

ISIS, School of Electronics and Computer Science,University of Southampton, Southampton, U.K.

`www.johndavidbustard.com`, `msn@ecs.soton.ac.uk`

Abstract

Recent work suggests that the human ear varies significantly between different subjects and can be used for identification. In principle, therefore, using ears in addition to the face within a recognition system could improve accuracy and robustness, particularly for non-frontal views. The paper describes work that investigates this hypothesis using an approach based on the construction of a 3D morphable model of the head and ear. One issue with creating a model that includes the ear is that existing training datasets contain noise and partial occlusion. Rather than exclude these regions manually, a classifier has been developed which automates this process. When combined with a robust registration algorithm the resulting system enables full head morphable models to be constructed efficiently using less constrained datasets. The algorithm has been evaluated using registration consistency, model coverage and minimalism metrics, which together demonstrate the accuracy of the approach. To make it easier to build on this work, the source code has been made available online.

1. Introduction

In the field of face recognition, morphable model fitting has been used very effectively to identify people under relatively unconstrained settings [6]. However, evaluations of these techniques show that there is still significant scope for improvement [18]. One possibility is to include additional recognition features. The ear is particularly suitable for this purpose as it has a wide variation in appearance between individuals and, like the face, is recognisable at a distance. It also has some advantages over the face in that its appearance does not alter with expressions, is rarely disguised by makeup or cosmetic surgery, and is believed to remain similar in appearance with age.

Earlier work has confirmed the ear as a viable feature for recognition using two dimensional techniques [10] [11] [1] [13]. However, results are sensitive to large pose or lighting changes so an alternative approach based on the construction of a morphable model of the face

and ear is now being investigated.

Existing morphable models of the head have focused on the face and implicitly or explicitly avoided accurate ear reconstruction [5] [2]. As a result, range scan data of the ear is generally of lower quality and less complete than that available for the face [21]. This neglect of the ear is partly due to the challenge of modelling its more detailed and self occluding structure. In addition, ears have not been a priority in existing work as they are not generally used by humans for recognition.

The main contribution of the work described here is a novel technique for the morphable model construction of a face profile and ear using noisy, partial and occluded data. The resulting system is the first developed for modelling the three dimensional space of ear shapes. The model is constructed by registering a generic head mesh with 160 range scans of face and ear profiles. Occluders and noise are identified within the scans using an automated classifier. The remaining valid regions are then used to register the mesh using a robust non-linear optimisation algorithm. Once registered, the scan orientations are normalised and then used to construct a linear model of all head shapes.

The next section summarises relevant existing work on morphable model construction and the representation of ears in those models. This is followed by a discussion of how the fitting problem can be formalised and the technique is then described in detail. Particular attention is given to the automated process for removing noise and occlusions in the training data. Finally, an evaluation section describes how three model metrics are measured and summarises the experimental results obtained. The paper concludes with proposals for future work. The algorithms used in this paper are available through the project website [9].

2. Related work

In 1999, Blanz and Vetter created the first 3D morphable model [7]. It was constructed from over 200 cylindrical range and colour scans of male and female heads, registered with each other using an optical flow algorithm. The model was constructed using the mean of these values and their first 90 eigenvectors calculated using PCA (principal com-

Figure 1. This image is of the base mesh, cleaned scan and fitted model for the technique used by Amberg et al. [3] It shows that the ear is not affected by the range scan and retains the shape of the base mesh
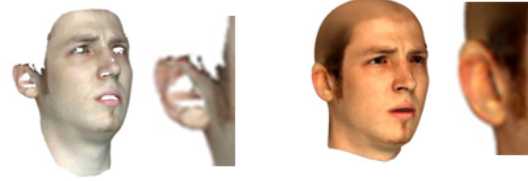

Figure 2. Results from the work by Basso et al. [5] showing the base mesh, range scans and the resulting registered models. Note that the registered models have less protruding ears than the original scans and lack internal detail
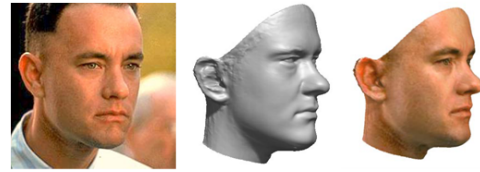

Figure 3. Image of a fitted Blanz and Vetter ear [7], the inner detail of which is missing but is compensated by the information in the texture
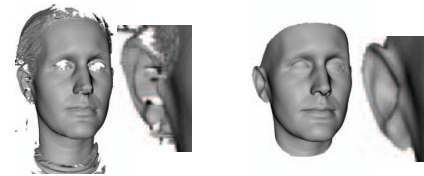

Figure 4. Image of a fitted Paysan et al. [14] ear, much of the ear is accurately modelled. However, the top left outer curve and middle ear hole covering are not fully reconstructed.

ponents analysis). This produced a highly realistic model of face appearances which they then used to create photorealistic 3D meshes from single photographs.

Further work has concentrated on improving the registration process. For example, in 2006, Basso et al. extended the optic flow approach to handle expression varying datasets [5]. In the same year, Vlasic et al. produced a multilinear morphable model that enabled independent adjustment of identity and expressions [20]. This used an optimisation-based fitting technique. Then in 2008 Amberg et al. [3] replaced the general optimisation framework with a non-rigid iterated closest point algorithm. With their approach they were able to construct models using partial range scans. Another contribution from Patel et al. was to demonstrate the importance of Procrustes normalisation of scans before calculating the subspace and they proposed an alternative fitting algorithm using thin-plate splines between manually labelled feature points [17].

Other researchers have applied the morphable model approach to different objects. For example in 2005, Allen et al. constructed morphable models of human body shapes [2]. In similar work, Anguelov et al. [4] developed a novel automated registration algorithm for bodies with significant pose variation. Their technique enabled the construction of a model to estimate a subject's body shape under different poses.

These existing approaches have concentrated on face or body shape and have generally, explicitly or implicitly avoided constructing accurate ear models. For example, in the work of Allen et al. the ear region of their training data is marked as inaccurate, which removes its influence on model construction. This produces heads on which all ears have the same shape. Similarly Amberg et al's and Basso et al's fitting results show little influence of the ear on the fitting process, as indicated in Figures 1 and 2.

Of the existing research, some of the most accurate ear models have been produced by the original Blanz and Vetter model (Fig. 3). Their work approximated the head by regularly sampling points over its surface and then connecting them to form a complete manifold. This technique creates ear models that lack folds and self occlusion but as the surface texture includes some shadow information the visual results look reasonably accurate.

In addition, a recent paper by Paysan et al. [14] has demonstrated high quality head models (Fig. 4). These were constructed using a more precise scanning process and registered using the algorithm of Amberg et al. This is in contrast to the work described in this paper, which focuses on achieving accurate results using less controlled datasets through the use of an automated occlusion and noise classifier. All of the above techniques involve some degree of manual labelling to achieve accurate results. However, early work by Brand [8] demonstrated the potential for automatically constructing morphable models directly from video sequences. This approach used optical flow, structure from motion, and novel matrix factorisation techniques. The resulting models are visually convincing but are only demonstrated on single individuals and lack the detail of models produced by registered range scans.

3. Defining the problem

One current difficulty in developing new types of morphable model is that there is no recognised definition for an

optimal model; that is, there is no agreement on the metrics to be used to optimise the model parameters.

Initial morphable model construction work used qualitative rather than quantitative evaluations of accuracy. In each case, the precision of the techniques were shown visually through a number of rendered example registrations [7]. One exception to this, however, is the work of Amberg et al. [3], which includes two measures of the quality of their created models. The first was based on the accuracy with which excluded training samples could be reconstructed using a model built from the remaining data. The second measure evaluated the average angle between the normals of corresponding points on separate registered models. This quality value was determined by calculating this difference between every possible pair of models in the training set and averaging the results. The rationale behind this measure was that correctly registered model parts should have correspondingly similar surface normals.

This error measure can be formalised as

$$error(i,j) = \frac{1}{l} \sum_{k=0}^{l} cos^{-1}(vn(i,k), vn(j,k))^2$$

Where i and j are the two aligned models, $vn(i,k)$ is a function that returns the normal of the vertex k of the model i, and l is the number of vertices in the model.

A similar evaluation was provided by Patel et al. [17] where they projected labelled points on range scans not included in the training set and then estimated the resulting reconstruction error when the models had been fitted to the images. In this way the value of the model for accurate reconstruction from images could be estimated.

To build on this work the paper explores two desirable morphable model properties: *coverage* and *minimalism*. Model coverage refers to the degree to which the model can generalise to samples beyond its training set and model minimalism refers to the amount of redundancy in the model; in essence, the size of non-head space that is covered. Unfortunately, calculating these properties directly is ill-posed and computationally expensive. This is due to the extremely large number of parameters involved and the relatively small quantity of training data available. In the case of morphable models of a class of object, such as all heads, existing approaches use heuristic techniques and assumptions to make the problem tractable.

The primary underlying assumption of existing morphable model work is that by creating an accurate registration between a large set of head models a linear subspace can be created that accurately approximates all head shapes. A difficulty here, however, is that the registration is not unique, as there is no agreement on which parts of a modelled object should be considered the same between different subjects. Existing approaches have tackled this issue in different ways.

With the exception of the automated technique of Brand [8], all current techniques initialise their registrations using a set of manually labelled features, followed by the application of an algorithm that calculates a dense correspondence. In the work of Patel et al. [17] a thin plate spline is used to interpolate a cylindical mapping calculated from the manually labelled points. In contrast, Vlasic et al. [20] use an iterative approach based on a similar principle to the iterative closest point algorithm, progressively adjusting the fitting using closest points as an estimate of the correspondence between mesh and range data. This was refined by Amberg [3] to calculate the optimal least deformation solution to the correspondence at each iteration. Finally, both Blanz et al. [7] and Anguelov et al. [4] include local appearance similarity measures in their fitting techniques.

Blanz and Vetter's work formulated the problem as an optic flow calculation between cylindrical range scans. In contrast, Anguelov et al. calculated local mesh signatures using Spin images [15]. Spin images are a translation and rotationally invariant signature that describe the surface shape of a range scan. By using this feature their approach could be applied to less constrained range meshes. In addition, their correlated correspondence algorithm calculated the global minimum for both the local similarity and the local smoothness of each part of a mesh using loopy belief propagation. Using this technique they were able to construct accurate full body morphable models with significant variation in body pose.

Overall, the accuracy and consistency of these techniques is dependent on the quality of the initial labelling and the smoothness of their cost functions. Their relative performance can be evaluated by analysing the resulting morphable models or through direct metrics that evaluate each fitting individually. In the work described in this paper, the quality of each registration has been evaluated using the consistency of the resulting mesh. The consistency of the fitting is calculated by comparing the similarity of meshes produced when multiple range images of the same person are fitted independently. If multiple scans of the same person result in meshes that are more similar than those of any of the other subjects then it is likely that the fitting process and the resulting model can be used for recognition purposes.

In addition, normalisation can be applied to improve the minimalism in the constructed model. Normalisation involves aligning registered head models with one another so they can be expressed by a minimal subspace. This is based on the observation that the minimal subspace removes variation due to differing head poses, rather than differing surface shapes. This is most obvious when a number of identical meshes at different poses are used to construct a model. The minimal subspace is then achieved when all the models are perfectly aligned.

The technique described in this paper is based on the approach developed by Vlasic et al. [20]. By using this general non-linear optimisation framework, constraints can easily be adapted and incorporated without a major change in the underlying algorithm. The details of this technique and its evaluation are outlined in the next section.

4. Technique

4.1. Training data

Where possible, it is desirable to use existing datasets as a source of representative 3D scans. One of the largest datasets is that provided for the Face Recognition Grand Challenge [18]. This data is in the form of front facing range data and colour values. The range data is estimated to be accurate to within 0.4 mm. As the images are front facing, however, the ear data detail is restricted. A smaller, but more relevant source is the Notre Dame Biometric Database J [21]. This dataset was obtained with the same sensor as the FRGC data and was designed specifically for ear recognition. It contains a number of profile head images and has been used to evaluate various existing 3D ear recognition algorithms.

The original work by Blanz and Vetter used complete head scans with minimal noise and uniform lighting. As a result, they could perform an optic flow calculation based on the surface appearance. However, the Notre Dame data is less constrained and less complete. In particular, it does not cover the entire head surface and is recorded with varying poses and lighting conditions. This makes an optic flow calculation less reliable. More recent work, such as that of Allen et al. [2], has instead focused on deforming a generic base mesh to align it with the range data. The aligned mesh parameterises the range data and registers the samples. This is the framework that has been used in the approach described here.

4.2. Preprocessing

The scans used by Blanz and Vetter were preprocessed manually to remove hair and provide an initial head alignment. They also made their subjects wear caps to minimise hair occlusion. Allen et al. also used caps [2]. In addition, they marked regions of the base mesh as being scanned inaccurately. These points then had no influence on the fitting process. The range scan data used in this paper only covers a part of the head volume. In addition, the ear regions contain significant self-occlusion. General hole filling algorithms have been developed that could complete these scans [12]. However, the internal complexity of the ear is likely to produce incorrect results as these algorithms make smoothness and convexity assumptions that are not valid for an ear shape. For this reason, the mesh registration and model construction steps have been designed to work

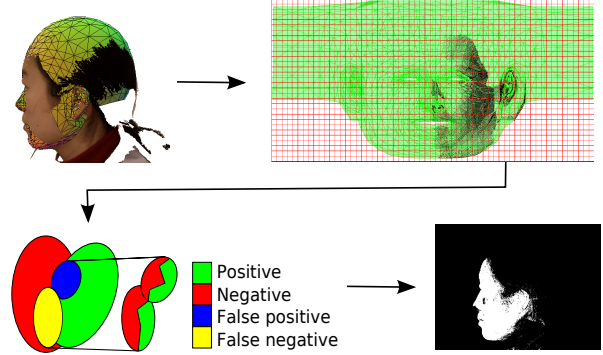Figure 5. Left: A range scan of the ear Right: A tilted view revealing inaccurate mesh regions

Figure 6. A flow chart showing the classification algorithm

with partial data. The Notre Dame data contain partial occlusions as well as noise generated by the range scanning process. This is particularly noticeable near the Intertragic Notch where there is a sudden change in depth. This region is circled within Figure 5.

To address these issues, the model construction process has been adapted to be robust to both occlusion and noise. This is achieved by labelling parts of the mesh as invalid. These invalid regions then do not contribute to the fitting process. The labelling process is performed automatically using a classifier trained on 30 manually labelled images. Figure 6 shows a flow chart summarising this process.

To classify the surface, the model was first aligned and deformed using a set of labelled feature points. The range scan was then given an approximate registration by calculating the closest mesh surface points to range scan pixels. Using this registration each point on the mesh surface was assigned a two dimensional uv location. This uvmap of points was then split into a number of overlapping regions. Each region was used to train a separate classifier using the Spin images of each pixel of the manually labelled examples. The classifier approximates the space of positive and negative samples with a pair of Gaussian models. These models are then used to estimate the class of each training sample, with the closer model being used as an estimate of validity. The values that are incorrectly estimated by this process are grouped into sets of false positive and false negative
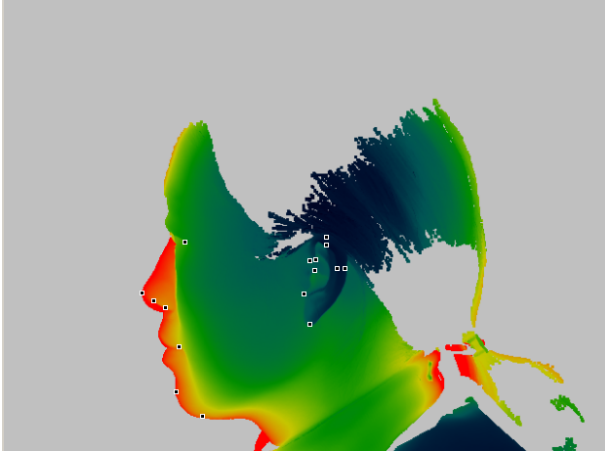
Figure 7. The location of the manually identified feature points

Figure 8. Image of least squared rigid registration of base mesh to feature points

samples. Each of these samples is used to calculate Gaussian models representing their regions. Within each region memory efficient kdtrees are constructed for classification. These kdtrees are then used to evaluate the k nearest neighbours' estimate of each region. In this way the majority of potential range samples are classified with the efficiency of a simple Gaussian classifier whilst still maintaining the accuracy of a detailed decision boundary between classes. Once classified, each overlapping region contributes votes to the validity of its pixels. The estimate with the most votes is returned as the classification. In cases where the votes are equal a negative classification is returned.

The classification accuracy is analysed in the evaluation section. This classifier is straightforward to implement and has the added advantage that it can be constructed efficiently for large datasets.

4.3. Feature points

Similar to existing work, the base mesh is initialised using manually placed feature points, as indicated in Figure 7. These points were selected through experimentation to address visible errors in the fitting of a training sample set whilst minimising the number of required feature points.

4.4. Initial registration

Using the 3D position of the marked feature points, the Procrustes algorithm is used to calculate the least-squared error rigid transformation between the base mesh feature points and those in the range scan (Fig. 8). Once initialised a more precise fit is obtained.

4.5. Optimisation based fitting

The fitting optimisation problem can be formulated in a number of different ways. In all cases, however, the goal is

to define the optimisation such that when the error values are minimised the resulting registration will lead to an optimal model. In addition, the ideal error values are those that smoothly and monotonically decrease towards a solution. Under these conditions, existing non-linear least squares based optimisations can be used to find a solution. Also, to obtain these results efficiently, it is desirable for the problem to be expressed with parameters that each influence a small number of error values. This will result in zero values for most elements of the matrices of the optimisation problem. As a result, efficient sparse variants of linear algebra methods can be used.

The challenge with morphable model construction is that it is not obvious how best to define the problem so that a minimal value corresponds to an optimal result. However, consistent with existing work, the technique described here minimises the following properties:

Distance to Feature Points: The features points are labelled on the range image by selecting the range point that most closely represents the feature. The error is a three dimensional value representing the relative displacements of the Mesh features in the X, Y and Z directions. The feature points are defined on the mesh as points on the triangles of the surface of the mesh. The movement of the vertices of the triangle are then the only parameters that influence these error values.

Smooth Deformation: In the work of Allen et al. [2] and Vlasic et al. [20] this error is formulated by associating a homogeneous transformation matrix with each vertex of the mesh to be aligned. Smoothness is then maintained by minimising the Frobenius norm between the matrices of connected vertices, where connection is determined by the edges of the mesh. This causes a preference for regions of consistent local affine transformation, namely rotation, scaling and shearing. These operations maintain the surface continuity and preserve local detail, resulting in a smooth deformation. In this way the parameters of the fitting process are the elements of the matrices of each vertex.

Alignment of Mesh Surface to Range Points: Each valid range point has the closest mesh surface point estimated. These are parameterised using barycentric coordinates of the closest triangle. Each point places a constraint on the

three vertices defining the triangle. This is in contrast to existing work which uses the distance between mesh vertices and their closest range points as an error term. The approach proposed here enables a more accurate alignment of any given resolution of mesh.

Optimisation Algorithm: Using the constraints defined in the preceding paragraphs the optimisation algorithm can be constructed. The constraints have a large number of parameters. In addition, they have a non-linear effect on their error values. These constraints can be solved using existing non-linear optimisation algorithms. For general smooth problems, such as these, one of the most efficient is the Levenberg-Marquart optimisation algorithm [16] formulated to take advantage of sparse constraints.

At each iteration of the Levenberg-Marquart optimisation, the error values associated with the distance of the mesh surface to the range points are sorted and the largest 1% of the errors are estimated to be false positives in the surface classification process. These errors and their associated constraints are then excluded from the optimisation. The percentage of constraints excluded has been estimated manually to achieve the best balance between excluding misclassified surface constraints and using the maximum amount of information to register the mesh accurately.

4.6. Data normalisation

Once a number of range images have been registered, the model can be created. This involves calculating the mean and principal components of the shape. However, the range meshes may contain variations due to the pose of the subjects when they were recorded. These variations are partially addressed through the initialisation using feature points but may not represent the optimal normalisation necessary to construct a minimal model. To address this problem, models are normalised using the Procrustes transform.

An added complication is that fitted meshes are only valid for part of the surface due to noise and occlusion of the range data. In this case a valid subset of the data is used for normalisation. To address this issue each fitted mesh is registered to the original base model shape using the vertices that were constrained by the valid range data. The position of the remaining vertices are then estimated using the smoothness constraints in Section 4.5. This creates a smooth head model with no visible artefacts and exploits the implicit anatomical information contained within the base mesh. As all of the head samples cover a similar surface region this normalised head shape remains similar between multiple scans of the same subject. This allows a direct PCA algorithm to be used to construct the model. If this were not the case it would be necessary to use a PCA variant that operated with partial data, such as Probabilistic PCA [19].
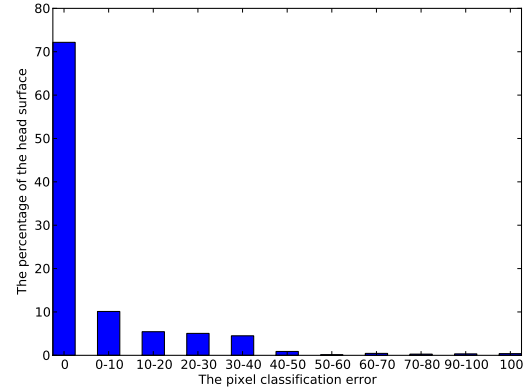

Figure 9. Graph showing the percentage of the head region that can be classified within a given error range

5. Evaluation

The work here builds on existing approaches by examining a number of metrics for determining the quality of the registration and the resulting model. The model has been constructed using 160 training images taken from the Notre Dame Biometric Database J [21].

5.1. Automated classification accuracy

To evaluate the accuracy of the noise and occlusion classifiers, ten-fold cross validation was applied to 30 hand-labelled training samples. Each region's accuracy was calculated using 10-fold cross validation. The distribution of accuracies can be seen in Figure 9 where the percentage of regions within a given error range is identified. Over half of the head surface is either not visible within the training set or consists exclusively of skin or occluder samples. The majority of the remaining regions have been correctly classified, but there is still a significant number of misclassifications. This is compensated for to some extent by the robust fitting process but still contributes error to the resulting model. It is likely that additional information, such as surface colour, will further improve these results.

5.2. Fitting consistency

The fitting consistency evaluates the uniqueness of each registration by comparing the similarity of meshes produced when two range images of the same person are fitted independently. This measure is calculated as the average distance between vertices of triangles that have been constrained by both range images. To compensate for variations in the pose of heads in each scan, the meshes are normalised to the base mesh using the Procrustes transform applied to the shared vertices. Figure 10 shows the relative distances between the closest registered head and the next
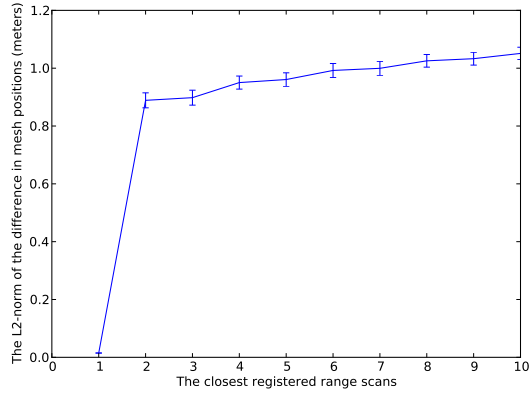
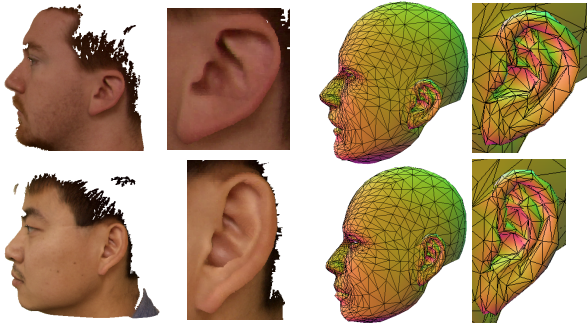Figure 10. Graph of fitting consistency

Figure 12. Graph of coverage

Figure 11. Two examples of registered head models

Figure 13. Graph of minimalism

closest nine scans. For all 160 samples within the training set, the closest head is the scan of the same person taken at a later date. The significant difference between the closest and the other scans indicates that the registration process is effective in extracting a consistent shape that can be used for recognition. Examples of these registrations can be seen in Figure 11.

5.3. Model metrics

The model has been evaluated using the criteria of coverage and minimalism as defined in Section 3. The results are:

Coverage: this is measured using a ten-fold cross validation technique, implemented by excluding a set of training images from the model construction process and then measuring the accuracy with which the fitted images can be reconstructed using the model. This is evaluated by projecting the fitted meshes into the space of the model and measuring the L2-norm of the difference in vertex positions between the fitted meshes and their representation using the model. The results can be seen in Figure 12.

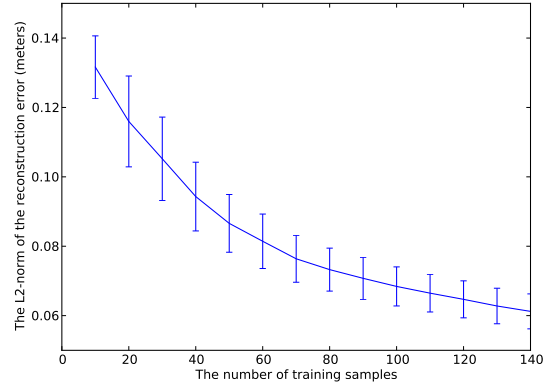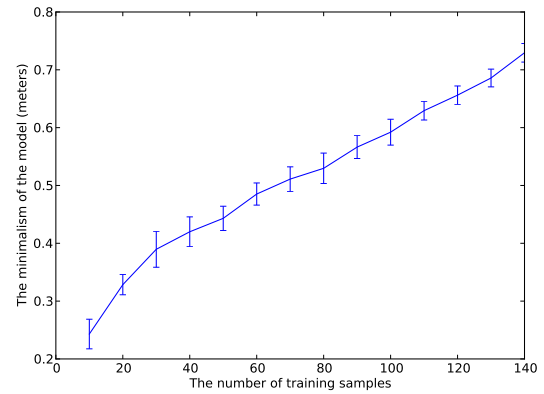*Minimalism*: this has been estimated using the square

root of the sum of the square of the eigenvalues of each model. This metric can be formalised as

$$minimalism(\mathbf{e}) = \sqrt{\sum_{k=0}^{l} e_k^2}$$

where $\mathbf{e}$ represents the model's eigenvalues.

As with coverage, the error in these values has been calculated using ten-fold cross validation. Figure 13 shows these results.

The graph of coverage shows a much faster rate of convergence than that for minimalism. This can be explained by errors within the registration process increasing the minimalism of the model without significantly improving its coverage.

6. Conclusions

This work demonstrates the first complete head morphable model explicitly designed to recreate both the face

and ear shape accurately. In addition, it improves the robustness of existing morphable model construction techniques by using classifiers trained to detect occluding and high noise areas. It also provides a framework for the evaluation of morphable models and uses this framework to highlight the advantages and sensitivities of the proposed technique. The evaluation shows that the described technique is extracting a consistent shape associated with the identity of the individual and that within the error margins of the registration process the 160 training samples are close to achieving convergence of the model.

Further work in this area will involve the use of the evaluation framework to examine other morphable model techniques, such as the correlated correspondence algorithm of Anguelov et al. [4] or the non-rigid iterated closest point algorithm developed by Amberg et al. [3] In addition, this framework can be expanded to evaluate the utility of the model in inferring identity from partial data such as detected features within 2D images.

There is also scope for creating an improved surface classifier, through additional distance constraints, and providing improved speed and precision using support vector machines or boosting techniques. If this could be combined with an automatic feature point detection process, fully automated model construction from relatively unconstrained training data would be possible. This offers the potential for more widespread application of morphable models within object recognition.

References

[1] M. Abdel-Mottaleb and J. Zhou. Human ear recognition from face profile images. In *ICB*, pages 786–792, 2006.

[2] B. Allen, B. Curless, and Z. Popović. The space of human body shapes: reconstruction and parameterization from range scans. In *SIGGRAPH '03: ACM SIGGRAPH 2003 Papers*, pages 587–594, New York, NY, USA, 2003. ACM.

[3] B. Amberg, S. Romdhani, and T. Vetter. Optimal step non-rigid icp algorithms for surface registration. In *Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on*, pages 1–8, June 2007.

[4] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. Scape: shape completion and animation of people. *ACM Trans. Graph.*, 24(3):408–416, 2005.

[5] C. Basso, P. Paysan, and T. Vetter. Registration of expressions data using a 3d morphable model. In *FGR '06: Proceedings of the 7th International Conference on Automatic Face and Gesture Recognition*, pages 205–210, Washington, DC, USA, 2006. IEEE Computer Society.

[6] V. Blanz, S. Romdhani, and T. Vetter. Face identification across different poses and illuminations with a 3d morphable model. In *Automatic Face and Gesture Recognition, 2002. Proceedings. Fifth IEEE International Conference on*, pages 192–197, May 2002.

[7] V. Blanz and T. Vetter. A morphable model for the synthesis of 3d faces. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.

[8] W. Brand. Morphable 3d models from video. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, volume 2, pages II–456–II–463 vol.2, 2001.

[9] J. D. Bustard. 3d morphable model construction for robust ear and face recognition, May 2010. www.johndavidbustard.com/mmconst.html.

[10] J. D. Bustard and M. S. Nixon. Robust 2d ear registration and recognition based on sift point matching. In *Biometrics: Theory, Applications and Systems, 2008. BTAS 2008. 2nd IEEE International Conference on*, pages 1–6, 29 2008-Oct. 1 2008.

[11] K. Chang, K. Bowyer, S. Sarkar, and B. Victor. Comparison and combination of ear and face images in appearance-based biometrics. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 25(9):1160–1165, Sept. 2003.

[12] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. In *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 71–78, New York, NY, USA, 1992. ACM.

[13] D. J. Hurley, B. Arbab-zavar, and M. S. Nixon. The ear as a biometric. In *Handbook of Biometrics. Forthcoming*, 2007.

[14] IEEE. *A 3D Face Model for Pose and Illumination Invariant Face Recognition*, Genova, Italy, 2009.

[15] A. Johnson. *Spin-Images: A Representation for 3-D Surface Matching*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, August 1997.

[16] K. Levenberg. A method for the solution of certain nonlinear problems in least squares. *Quarterly Journal of Applied Mathmatics*, II(2):164–168, 1944.

[17] A. Patel and W. Smith. 3d morphable face models revisited. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, 0:1327–1334, 2009.

[18] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the face recognition grand challenge. In *CVPR '05: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Volume 1*, pages 947–954, Washington, DC, USA, 2005. IEEE Computer Society.

[19] M. E. Tipping and C. M. Bishop. Probabilistic principal component analysis. *Journal Of The Royal Statistical Society Series B*, 61(3):611–622, 1999.

[20] D. Vlasic, M. Brand, H. Pfister, and J. Popovic. Face transfer with multilinear models. In *ACM Transactions on Graphics (TOG)*, volume 24, pages 426–433, New York, NY, USA, July 2005. ACM.

[21] P. Yan and K. W. Bowyer. Biometric recognition using three-dimensional ear shape. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 29, pages 1297–1308, 2007.